arXiv:1405.4538v2 [stat.ME] 9 Jan 2016

Unit-free and robust detection of differential expression from RNA-Seq data

Hui Jiang[1,2,*] and Tianyu Zhan[1]

[1]Department of Biostatistics, University of Michigan
[2]Center for Computational Medicine and Bioinformatics, University of Michigan
[*]Please send correspondence to jianghui@umich.edu.

January 12, 2016

Abstract

Ultra high-throughput sequencing of transcriptomes (RNA-Seq) is a widely used method for quantifying gene expression levels due to its low cost, high accuracy and wide dynamic range for detection. However, the nature of RNA-Seq makes it nearly impossible to provide absolute measurements of transcript abundances. Several units or data summarization methods for transcript quantification have been proposed in the past to account for differences in transcript lengths and sequencing depths across different genes and different samples. Nevertheless, further between-sample normalization is still needed for reliable detection of differentially expressed genes. In this paper we propose a unified statistical model for joint detection of differential gene expression and between-sample normalization. Our method is independent of the unit in which gene expression levels are summarized. We also introduce an efficient algorithm for model fitting. Due to the L0-penalized likelihood used in our model, it is able to reliably normalize the data and detect differential gene expression in some cases when more than 50% of the genes are differentially expressed in an asymmetric manner. Comparisons with existing methods is given.

1 Introduction

Ultra high-throughput sequencing of transcriptomes (RNA-Seq) is a widely used method for quantifying gene expression levels due to its low cost, high accuracy and wide dynamic range for detection (Mortazavi et al., 2008). As of today, modern ultra high-throughput sequencing platforms can generate hundreds of millions of sequencing reads from each biological sample in a single day. RNA-Seq also facilitates the detection of novel transcripts (Trapnell et al., 2010) and the quantification of transcripts on isoform level (Jiang and Wong, 2009; Salzman et al., 2011). For these reasons, RNA-Seq has become the method of choice for assaying transcriptomes (Wang et al., 2009).

In a typical RNA-Seq experiment, mRNA transcripts are extracted from biological samples, reverse transcribed into cDNA molecules, randomly fragmented into pieces, filtered based on fragment lengths (size selection), linked with sequencing adapters and finally processed by a sequencer. The data output by the sequencer are sequenced reads (or reads for short), from either one end (i.e., single-end sequencing) or both ends (i.e., paired-end sequencing) of the fragments. These reads are usually aligned to reference transcripts or genomes and the data are then summarized as read counts for each transcript in each sample. Complications may occur when reads cannot be uniquely aligned to reference transcripts or genomes, either due to sequence homology among genes, or due to multiple transcripts (isoforms) sharing commons regions (exons) within a single gene. In this paper, we will ignore these complications and assume that each gene has only one transcript and therefore we will use terms "gene" and "transcript" interchangeably. We will also assume that each read can be uniquely aligned to a single gene, and we take read counts for each gene in each sample as our data. However, our approach can work with estimated read counts or gene expression levels from methods developed to handle these complications such as Jiang and Wong (2009) and Li et al. (2010). An overview of these methods is given in Pachter (2011).

One major limitation of RNA-Seq is that it only provides relative measurements of transcript abundances. Because reads are sequenced from a random sample of transcript fragments, changing the total amount of transcripts in a sample will have little effect on the distribution of finally sequenced reads. Furthermore, longer transcripts will generate more fragments which will subsequently result in more reads, and sequencing with higher depth will also lead to more reads for every transcript. To account for these factors,

several data summarization units (a.k.a. within-sample normalization methods) for transcript quantification have been proposed in the past to account for different sequencing depths across samples and possibly also for different transcript lengths across genes, which include CPM/RPM (counts/reads per million) (Robinson and Oshlack, 2010), RPKM/FPKM (reads/fragments per kilobase of exon per million mapped reads) (Mortazavi et al., 2008; Trapnell et al., 2010) and TPM (transcript per million) (Li et al., 2010). Since a "read" can refer to either a single-end read or a paired-end read, depending on the sequencing experiment conducted, we will use the term RPKM to represent both RPKM and FPKM in this paper. Suppose that there are a total of m genes in the sample. For $i = 1, \ldots, m$, let l_i be the length (often measured as the effective length after adjusted for edge effects) of gene i and let c_i be the observed read count for gene i in the sample. CPM (denoted as cpm_i), RPKM (denoted as $rpkm_i$) and TPM (denoted as tpm_i) values for gene i are defined as follows, respectively

$$\begin{aligned} cpm_i &= 10^6 c_i / \textstyle\sum_i c_i \\ rpkm_i &= 10^3 cpm_i / l_i \\ tpm_i &= 10^6 rpkm_i / \textstyle\sum_i rpkm_i \end{aligned} \tag{1}$$

In the past, all these methods have been used (with RPKM being the most widely used one) for quantifying gene expression levels from RNA-Seq data. It has been argued that TPM should be used instead of RPKM since TPM estimates the relative molar concentration of transcripts in a sample (Wagner et al., 2012). However, none of these methods can be used directly to detect differentially expressed (DE) genes reliably without a further between-sample normalization step, which is necessary to make gene expression measurements comparable across samples. Different between-sample normalization methods make different assumptions on the distribution of gene expression levels across samples. For instance, quantile normalization (Bolstad et al., 2003) assumes that the overall distributions of gene expression levels are the same for all the samples. In this sense, both RPKM and TPM can be considered as between-sample normalization methods when they are used directly to detect DE genes without additional normalization – TPM assumes that the total numbers of transcripts (i.e., total molar amount) are the same for all the samples and RPKM assumes that the total numbers of nucleotides in all the transcripts (i.e., total physical mass) are the same for all the samples, both of which are strong but arguably reasonable assumptions.

To understand the limitation of between-sample normalization, consider a simple hypothetical example of comparing gene expression profiles of two samples A and B. If 60% of the genes were up-regulated by 2-fold in sample B while the other 40% of the genes stayed stable in both samples, due to the relative nature of RNA-Seq measurements, it is impossible to distinguish it from the scenario where the first group of 60% genes actually stayed stable but the second group of 40% gene were down-regulated by 2-fold in sample B. Since the problem is non-identifiable in such cases, the normalization approach has to rule out the ambiguities based on its assumptions. One commonly used assumption is that the majority (i.e., $> 50\%$) of the genes are non-DE. The median-based approach (Anders and Huber, 2010) and TMM (trimmed mean of M values) (Robinson and Oshlack, 2010) are two normalization methods based on this assumption. Furthermore, between-sample normalization and detection of DE gene are two problems that are always tangled together, since ideally normalization should be based on non-DE genes only. The iterative normalization approach (Li et al., 2012) utilizes this idea and iterates between normalization and detection of DE genes. However, it is unclear what objective function is optimized in such an approach, and the final solution often depends on the initial guess which is undesirable. For an overview of between-sample normalization approaches and comparison of their performance for detection of DE genes, please refer to Dillies et al. (2013); Rapaport et al. (2013).

Fortunately, the example described above is rather extreme and unrealistic, because in practice genes rarely change at the same pace – it is unlikely that all the DE genes are up-regulated by the same amount (2-fold in the above example). Therefore, a practical assumption is that among DE genes, the degrees at which genes change have an unknown but spreaded distribution. This assumption, while largely ignored by most existing approaches, will be exploited in this paper.

In this paper, we will propose a unified statistical model for joint detection of differential gene expression and between-sample normalization. We will introduce the model and an efficient algorithm for model fitting in Section 2. Comparisons with existing methods will be given in Section 3, followed by discussions in Section 4.

2 A penalized likelihood approach

2.1 The model

Suppose there are a total of m genes measured in S groups of experiments with $n_1, \ldots, n_S$ samples, respectively. Let $x_{sij}, s = 1, \ldots, S, i = 1, \ldots, m, j = 1, \ldots, n_s$ be log-transformed gene expression measurements (a small positive number is usually added before taking logarithm) for the i-th gene from the j-th sample in the s-th group. The following statistical model is assumed

$$x_{sij} \sim N(\mu_{si} + d_{sj}, \sigma_i^2) \tag{2}$$

where μ_{si} is the mean of log-transformed expression levels of gene i in group s, d_{sj} is a scaling factor (e.g., log(sequencing depth) or log(library size)) for sample j in group s and σ_i^2 is the variance of log-transformed expression levels of gene i across all S groups. Here we assume that the log-transformation stabilizes the variances and makes them roughly the same across groups (yet can still be different across genes). Nevertheless, our model can be extended to accommodate heteroscedastic variances across groups.

Our main interest is in detecting differentially expressed (DE) genes across the S groups. Let τ_i be the indicator of differential expression for gene i such that $\tau_i = 1$ if gene i is differentially expressed across the S groups and $\tau_i = 0$ otherwise, i.e., $\tau_i = 1(\mu_{1i} = \mu_{2i} = \cdots = \mu_{Si})$ where $1(\cdot)$ is the indicator function. The parameters of major interest are $\{\tau_i\}_{i=1}^m$, while μ_{si}, d_{sj} and σ_i^2 might be of interest too, because they denote biologically meaningful quantities.

It can be easily shown that due to the log-transformation on the data, our method is independent of the unit (i.e., counts, CPM/RPM, RPKM/FPKM or TPM) in which gene expression levels are summarized, which is a desirable property in practice.

2.2 Penalized likelihood

For now, we assume that $\{\sigma_i^2\}_{i=1}^m$ are known. In practice, we solve for σ_i^2 using an ad hoc approach, which will be described later.

To fit model (2), we reparametrize μ_{si} as $\mu_i = \mu_{1i}, \gamma_{si} = \mu_{si} - \mu_{1i}, s = 2, \ldots, S$. Then the model becomes

$$\begin{cases} x_{1ij} \sim N(\mu_i + d_{1j}, \sigma_i^2) \\ x_{sij} \sim N(\mu_i + \gamma_{si} + d_{sj}, \sigma_i^2), s = 2, \ldots, S \end{cases} \tag{3}$$

To fit model (3), we minimize its negative log-likelihood

$$l(\mu,\gamma,d;x)=\sum_{i=1}^{m}\frac{1}{2\sigma_i^2}\left(\sum_{j=1}^{n_1}(x_{1ij}-\mu_i-d_{1j})^2+\sum_{s=2}^{S}\sum_{j=1}^{n_s}(x_{sij}-\mu_i-\gamma_{si}-d_{sj})^2\right)$$

where the term $(n/2)\sum_{i=1}^{m}\log(2\pi\sigma_i^2)$ is discarded as it does not contain any unknown parameter when σ_i^2 is known. Model (3) is non-identifiable because we can simply add any constant to all the d_{sj}'s and subtract the same constant from all the μ_i's, while having the same fit for $l(\cdot)$. To resolve this issue, we fix $d_{11}=0$. Furthermore, we introduce a penalty $p(\gamma)$ on all the γ_{si}'s and formulate a penalized likelihood

$$f(\mu,\gamma,d)=l(\mu,\gamma,d;x)+p(\gamma) \tag{4}$$

Commonly used penalty functions for $p(\gamma)$ are L1 (a.k.a. lasso) (Tibshirani, 1996), L0, SCAD (Fan and Li, 2001) and etc. The penalty function will force some γ_{si}'s to become zero, which will in turn facilitate the detection of DE genes since by definition $\tau_i=1(\mu_{1i}=\mu_{2i}=\cdots=\mu_{Si})=1(\sum_{s=2}^{S}|\gamma_{si}|>0)$.

Let $n=\sum_{s=1}^{S}n_s$ be the total sample size. There are mn observations and $mS+n-1$ free parameters in the model, and typically we can have n in tens or hundreds and m in tens of thousands. Using an L1 penalty, it will be computationally intensive if we fit the model using a lasso solver such as Glmnet (Friedman et al., 2010), and typically it will be computationally even more challenging to fit the model with a non-convex penalty such as L0 or SCAD. Fortunately, we can take advantage of the structure in model (3) and solve it efficiently. In this paper we work with the L0 penalty due to its robustness in estimation and variable selection

$$p(\gamma)=\sum_{i=1}^{m}\alpha_i 1(\sum_{s=2}^{S}|\gamma_{si}|>0) \tag{5}$$

where $\alpha_i>0, i=1,\ldots,m$ are tuning parameters. The approach to choose $\{\alpha_i\}_{i=1}^{m}$ will be described later. Our model fitting approach can also be adapted to accommodate other penalty functions.

2.3 Model fitting

It can be shown that the solution to (4) with penalty (5) can be obtained as follows (see Appendix for details).

Proposition 2.1. *Model (4) with penalty (5) can be solved as follows*

$$
\begin{aligned}
&d'_{sj} = (\textstyle\sum_{i=1}^m (x_{sij} - x_{si1})/\sigma_i^2)/(\sum_{i=1}^m 1/\sigma_i^2), s = 1, \ldots, S \\
&\mu'_{si} = (1/n_s) \textstyle\sum_{j=1}^{n_s} (x_{sij} - d'_{sj}), s = 1, \ldots, S \\
&d_1 = 0 \\
&d_2, \ldots, d_S = \underset{d_2, \ldots, d_S}{\arg\min} \sum_{i=1}^m \min\left(g(d_2, \ldots, d_S), \alpha_i\right) \\
&\textit{where } g(d_2, \ldots, d_S) = \frac{1}{2\sigma_i^2} \left\{ \sum_{s=1}^S n_s (\mu'_{si} - d_s)^2 - \frac{1}{n} \left[\sum_{s=1}^S (n_s(\mu'_{si} - d_s)) \right]^2 \right\} \\
&d_{sj} = d_s + d'_{sj}, s = 1, \ldots, S \\
&\gamma_{si} = \begin{cases} 0 & \textit{if } g(d_2, \ldots, d_S) < \alpha_i \\ \mu'_{si} - \mu'_{1i} - d_s & \textit{otherswise} \end{cases} \\
&\mu_i = \begin{cases} (1/n) \textstyle\sum_{s=1}^S n_s (\mu'_{si} - d_s) & \textit{if } g(d_2, \ldots, d_S) < \alpha_i \\ \mu'_{1i} & \textit{otherwise} \end{cases}
\end{aligned}
$$

The only computationally intensive step in Proposition 2.1 is to solve for $d_2, \ldots, d_S$, for which the following function is minimized

$$G(d_2, \ldots, d_S) = \sum_{i=1}^m \min\left(g(d_2, \ldots, d_S), \alpha_i\right).$$

Typically, function $G(\cdot)$ is non-convex and non-differentiable, and therefore very difficult to minimize. However, in low-dimensional cases, function $G(\cdot)$ can be minimized efficiently using exhaustive search. Examples with $S = 2$ and $S = 3$ are given in Figure 1 (where $G'(\cdot)$, a variation of $G(\cdot)$ is shown; see Section 2.5 for details) and Figure 2, respectively.

2.4 Choosing α_i

Choosing the tuning parameters in a penalized regression model is usually very tricky because it involves the bias-variance tradeoff (Hastie et al., 2009). In practice, cross-validation can often achieve reasonable performance, at the cost of additional computation and less robustness. Fortunately, there is a simple way to choose the tuning parameters in our model, which is based on the property of the solution in Proposition 2.1. Let $y_{sij} = x_{sij} - d_{sj}$ denote the normalized data. Then

$$\mu'_{si} - d_s = \frac{1}{n_s} \sum_{j=1}^{n_s} (x_{sij} - d'_{sj} - d_s) = \frac{1}{n_s} \sum_{j=1}^{n_s} y_{sij}$$

is the mean of y_{sij} for gene i in group s. The condition for $\gamma_{si} = 0$ in Proposition 2.1 can be rewritten as

$$\frac{1}{(S-1)\sigma_i^2}\left\{\sum_{s=1}^{S}\frac{\left(\sum_{j=1}^{n_s} y_{sij}\right)^2}{n_s} - \frac{1}{n}\left(\sum_{s=1}^{S}\sum_{j=1}^{n_s} y_{sij}\right)^2\right\} < \frac{2\alpha_i}{S-1}$$

whose LHS mimics the form of the F-statistic for one-way ANOVA models, which suggests we choose α_i as $((S-1)/2)F^*_{1-q}(S-1, n-S)$, where $F^*_{1-q}(\cdot)$ is the critical value for one-sided level q tests with the F-distribution. We set $q = 0.01$ in our experiments.

2.5 Simplification for two-group comparison

Here we study the simplest case of two-group comparison, which is the most widely used experimental design for differential gene expression study. For $S = 2$, the condition for $\gamma_{2i} = 0$ is

$$\frac{1}{2\sigma_i^2}\frac{n_1 n_2}{n}(\mu'_{2i} - \mu'_{1i} - d_2)^2 < \alpha_i$$

which can be rewritten as

$$|\mu'_{2i} - \mu'_{1i} - d| < \lambda_i$$

where

$$\lambda_i = \sqrt{\frac{2n\sigma_i^2\alpha_i}{n_1 n_2}}$$

is a tuning parameter that is alternative to and more convenient than α_i. Therefore, the solutions for d_2 (denoted as d for simplicity), γ_{2i} (denoted as γ_i for simplicity) and μ_i in Proposition 2.1 can be simplified as

$$d = \arg\min_d G'(d) = \arg\min_d \sum_{i=1}^{m}\frac{1}{\sigma_i^2}\min\left((\mu'_{2i} - \mu'_{1i} - d)^2, \lambda_i^2\right)$$

$$\gamma_i = \begin{cases} 0 & \text{if } |\mu'_{2i} - \mu'_{1i} - d| < \lambda_i \\ \mu'_{2i} - \mu'_{1i} - d & \text{otherswise} \end{cases}$$

$$\mu_i = \begin{cases} (n_1\mu'_{1i} + n_2(\mu'_{2i} - d))/n & \text{if } |\mu'_{2i} - \mu'_{1i} - d| < \lambda_i \\ \mu'_{1i} & \text{othersise} \end{cases}$$

Similarly, the condition for $\gamma_i = 0$ can be rewritten as

$$\frac{\left|(1/n_2)\sum_{j=1}^{n_2} y_{2ij} - (1/n_1)\sum_{j=1}^{n_1} y_{1ij}\right|}{\sigma_i\sqrt{1/n_1 + 1/n_2}} < \frac{\lambda_i}{\sigma_i\sqrt{1/n_1 + 1/n_2}}$$

whose LHS mimics the form of the t-statistic for two-sample comparison with pooled variance, which suggests we choose λ_i as $t^*_{1-\frac{q}{2}}(n_1+n_2-2)\sigma_i\sqrt{1/n_1 + 1/n_2}$, where $t^*_{1-\frac{q}{2}}(\cdot)$ is the critical value for two-sided level q tests with the t-distribution. The corresponding value for α_i is $(1/2)t^*_{1-\frac{q}{2}}(n_1+n_2-2)^2$ which is the same as $((S-1)/2)F^*_{1-q}(S-1, n-S)$ suggested in Section 2.4.

2.6 Solving for $\{\sigma_i^2\}_{i=1}^m$

To solve for $\{\sigma_i^2\}_{i=1}^m$, consider the negative log-likelihood function (with $\{\sigma_i^2\}_{i=1}^m$ being unknown parameters as well) for model (2) restricted to group s

$$l_s(\mu, d, \sigma^2; x) = \sum_{i=1}^{m}\left(\frac{n_s}{2}\log(2\pi\sigma_i^2) + \frac{1}{2\sigma_i^2}\sum_{j=1}^{n_s}(x_{sij} - \mu_{si} - d_{sj})^2\right)$$

Taking partial derivatives of $l_s(\cdot)$ with respect to μ_{si}, d_{sj} and σ_i^2 respectively and setting the partial derivatives to be zero, we have

$$\mu_{si} = \frac{1}{n_s}\sum_{j=1}^{n_s}(x_{sij} - d_{sj})$$

$$d_{sj} = \frac{\sum_{i=1}^{m}\frac{1}{\sigma_i^2}(x_{sij} - \mu_{si})}{\sum_{i=1}^{m}\frac{1}{\sigma_i^2}}$$

$$\sigma_i^2 = \frac{1}{n_s}\sum_{j=1}^{n_s}(x_{sij} - \mu_{si} - d_{sj})^2$$

Estimates for μ_{si}, d_{sj} and σ_i^2 can then be iteratively updated using the above three equations until converge, which is similar to the method of iteratively reweighted least squares for two-way ANOVA models with heteroscedastic errors. In our implementation, we use the following modified equation for updating σ_i^2 to reduce the estimation bias.

$$\sigma_i^2 = \frac{1}{n_s - 1}\sum_{j=1}^{n_s}(x_{sij} - \mu_{si} - d_{sj})^2$$

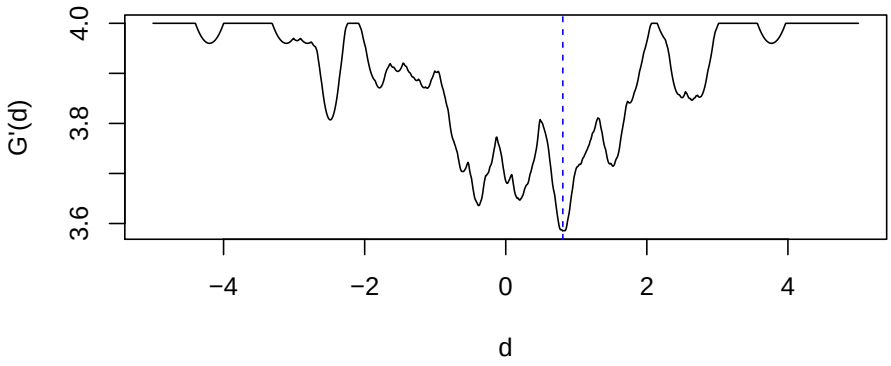

(a)

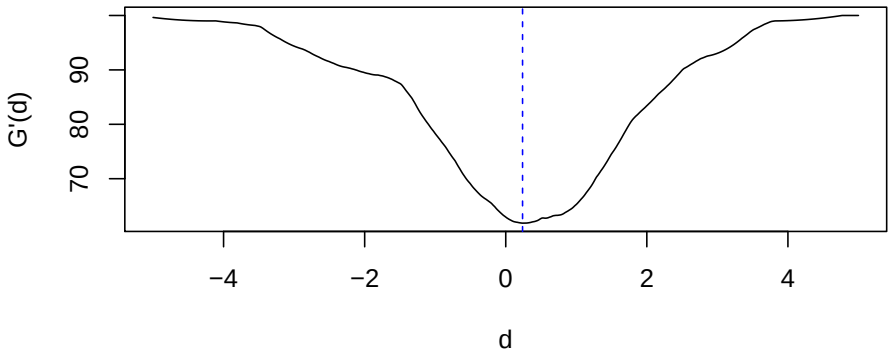

(b)

Figure 1: Function $G'(d) = \sum_{i=1}^{100}(1/\sigma_i^2)\min\left((\mu'_{2i} - \mu'_{1i} - d)^2, \lambda_i^2\right)$ with simulated $\{\mu'_{1i}\}_{i=1}^{100}, \{\mu'_{2i}\}_{i=1}^{100} \sim N(0,1), \{\sigma_i^2\}_{i=1}^{100} = 1$ with (a) $\{\lambda_i\}_{i=1}^{100} = 0.2$ and (b) $\{\lambda_i\}_{i=1}^{100} = 1$. The minimizers of $G'(\cdot)$ are shown with dashed lines.

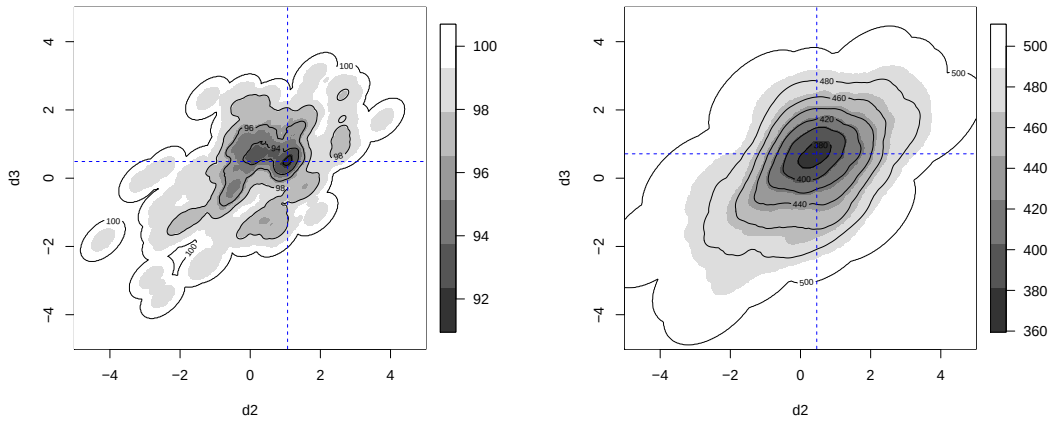

(a) (b)

Figure 2: Function $G(d2, d3)$ with simulated $\{\mu'_{1i}\}_{i=1}^{100}, \{\mu'_{2i}\}_{i=1}^{100}, \{\mu'_{3i}\}_{i=1}^{100} \sim N(0,1), \{\sigma_i^2\}_{i=1}^{100} = 1, n_1 = n_2 = n_3 = 10$ with (a) $\{\alpha_i\}_{i=1}^{100} = 1$ and (b) $\{\alpha_i\}_{i=1}^{100} = 5$. The minimizers of $G(\cdot)$ are at the intersections of the dashed lines.

Same as before, we fix $d_{s1} = 0$ and adjust the remaining d_{sj}'s and μ_{si}'s accordingly after each iteration to resolve the non-identifiability issue. Note that μ_{si} and d_{sj} estimated here are for group s only and therefore will be discarded. They should be be confused with the parameters estimated from model (2) based on the data from all the groups.

Denote the variance estimated using the above iterative algorithm for gene i in group s as s_{si}^2. We take a weighted average of $\{s_{si}^2\}_{s=1}^S$ to pool information from all the groups

$$s_i^2 = \frac{\sum_{s=1}^S (n_s - 1) s_{si}^2}{n - S}$$

Then we take another weighted average of s_i^2 and the estimated mean variance across all the genes to obtain a robust estimate for σ_i^2. That is

$$\widehat{\sigma_i^2} = (1 - w) s_i^2 + w \overline{s^2}$$

where $\overline{s^2} = \sum_{i=1}^m s_i^2 / m$, and the weight w is calculated using the following formula as suggested in Ji and Wong (2005) which is based on an empirical

Bayes approach

$$w = \frac{2(m-1)}{n-S+2}\left(\frac{1}{m} + \frac{(\overline{s^2})^2}{\sum_{i=1}^{m}(s_i^2 - \overline{s^2})^2}\right)$$

This kind of variance estimation approach is widely used in differential gene expression analysis with small sample sizes (Ji and Liu, 2010; Smyth, 2004). Using the estimated variances $\widehat{\sigma_i^2}$, we can then solve for d_s, γ_{si} and μ_i as described in Proposition 2.1.

3 Experiments

3.1 Simulation of two-group data

We simulate RNA-Seq data with a total of $m = 1000$ genes (700 non-DE genes, 300 DE genes) in two-group as follows

$\mu_{1i} = \mu_{2i} \sim N(-3, 2)$	mean log expression for genes 1-700
$\mu_{1i} \sim N(-3, 2), \mu_{2i} = \mu_{1i} + N(0, 1)$	mean log expression for genes 701-1000
$d_{sj} \sim N(0, 0.5)$	log scaling factor
$x_{sij} \sim N(\mu_{si} + d_{sj}, 0.2)$	log gene expression
$\log(l_i) \sim Unif(5, 10)$	log gene lengths
$N_{sj} \sim Unif(3, 5) \times 10^7$	library sizes
$\{c_{sij}\}_{i=1}^m \sim Mult\left(N_{sj}, \frac{l_i e^{x_{sij}}{}_{i=1}^m}{\sum_{i=1}^m l_i e^{x_{sij}}}\right) + 1$	read counts

We use $\log c_{sij}$ as data to fit our model. The fitted $\{\gamma_i\}_{i=1}^m$ are plotted in Figure 3a. Using CPM, RPKM or TPM values computed with formulas in (1) yield the same result. To demonstrate the robustness of our method, we also simulate with 300 non-DE genes and 700 DE genes. Furthermore, for DE genes we simulate with $\mu_{2i} = \mu_{1i} + N(1, 1)$, which means that the average log-fold change is 1. Our method still robustly estimates the γ_i's (Figure 3b). We further simulate with 100 non-DE genes and 900 DE genes, for which our method still achieves robust estimates when we simulate with $\mu_{2i} = \mu_{1i} + N(1, 1)$ (Figure 3c). Only when we simulate with 900 DE genes and with $\mu_{2i} = \mu_{1i} + N(3, 1)$, our method fails to achieve robust estimates (Figure 3d).

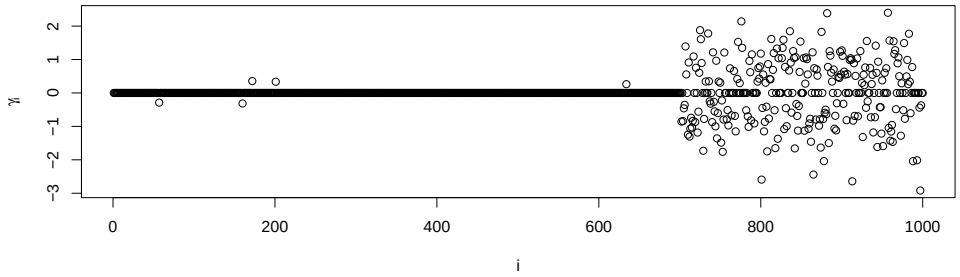

(a)

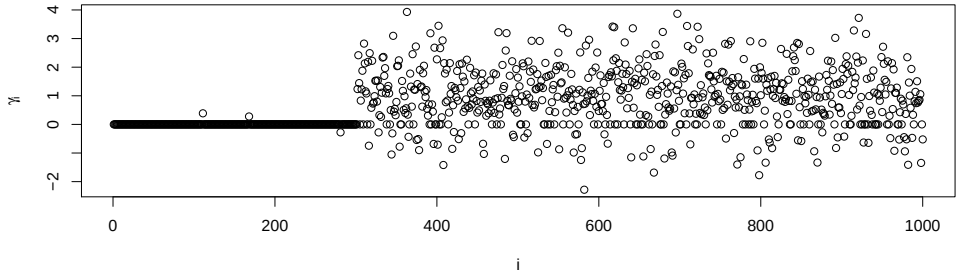

(b)

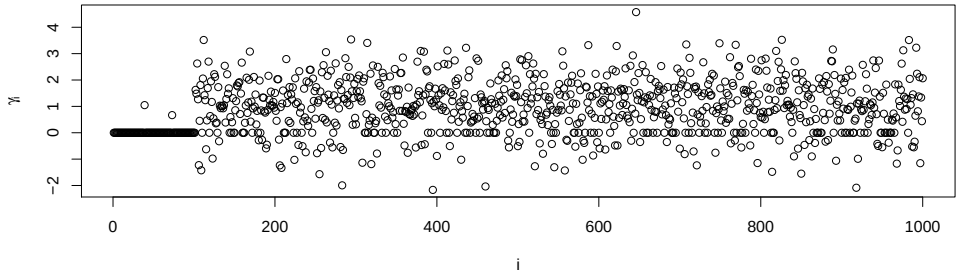

(c)

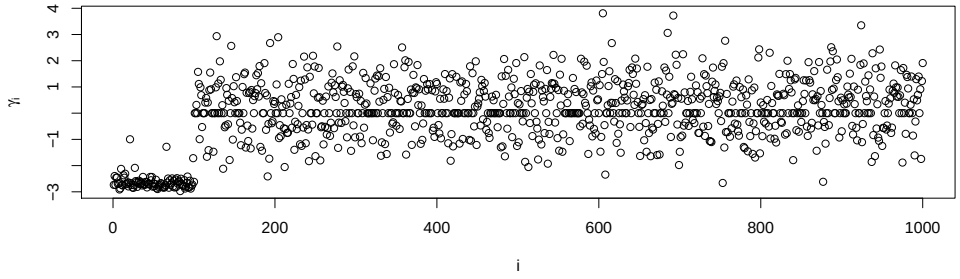

(d)

Figure 3: Estimated γ from simulated two-group data.

3.2 Comparison with existing methods

We compare our method (named rSeqRobust) with edgeR-robust (Robinson et al., 2010; Zhou et al., 2014), DESeq2 (Love et al., 2014), and limma-voom (Ritchie et al., 2015; Law et al., 2014), all of which are state-of-the-art methods for detecting differential gene expression from RNA-Seq data. We use two-group RNA-Seq data with a total of $m = 1000$ genes, and $n_1 = n_2 = 4$ samples in each group. We simulate both log-normally distributed read counts, which is the model assumptions of limma-voom and rSeqRobust, as well as negative-binomially distributed read counts, which is the underlying assumption of edgeR-robust and DESeq2. The distributions of gene expression levels and the distributions of library sizes for both simulations, as well as the distributions of read count dispersions for the negative-binomial simulation, are based on a real RNA-Seq dataset (Pickrell et al., 2010). In the log-normal simulation, log read counts are assumed to be normally distributed with $\sigma = 0.5$.

The read counts are simulated using the simulator described in (Zhou et al., 2014). We slightly modify the simulator to allow log-normally distributed data, as well as variable fold changes. We simulate data sets with 30% or 70% DE genes, as well as 50%, 70% or 90% up-regulated genes among all the DE genes. The log-fold change for DE genes (when measured as up-regulation from one group to the other) are assumed to be distributed as $N(\log 3, 1)$. We calculate AUCs for DE gene detection using all four methods, and summarize the results in Tables 1 and 2, respectively. For each parameter setting, the highest AUC value is shown in bold font. We can see that rSeqRobust is the best in log-normal simulations, and rSeqRobust and edgeR-robust are the best in negative-binomial simulations, which are expected. In particular, when rSeqRobust is outperformed by other methods, it is usually by only a small margin (i.e., with in one standard error) and only in relatively easier cases. In the most difficult cases, rSeqRobust outperform other methods by a relatively larger margin.

4 Discussion

It is shown in our simulations that our proposed approach is able to reliably normalize the data and detect differential expression in some cases when more than half of the genes are differentially expressed in an asymmetric manner.

Table 1: Comparison of edgeR-robust, DESeq2, limma-voom and rSeqRobust in log-normally distributed data. The table shows the percent of DE genes (DE %), percent of up-regulated genes among all the DE genes (Up %), as well as the mean AUCs for all four methods measured using 10 simulated replicates. The standard errors of the mean AUCs are given in parentheses.

DE (%)	Up (%)	edgeR-robust	DESeq2	limma-voom	rSeqRobust
30	50	0.8764	0.8684	0.8803	**0.8805**
		(0.0028)	(0.0024)	(0.0023)	**(0.0023)**
30	70	0.8547	0.8454	0.858	**0.8618**
		(0.0029)	(0.0038)	(0.0032)	**(0.0032)**
30	90	0.8446	0.8329	0.848	**0.8663**
		(0.0076)	(0.006)	(0.0079)	**(0.0076)**
70	50	0.872	0.8699	0.8775	**0.8776**
		(0.003)	(0.0025)	(0.0029)	**(0.0031)**
70	70	0.8151	0.8066	0.8195	**0.8434**
		(0.0067)	(0.0059)	(0.0074)	**(0.0064)**
70	90	0.6577	0.6336	0.6583	**0.7572**
		(0.0074)	(0.0074)	(0.0072)	**(0.0059)**

Table 2: Comparisons of edgeR, DESeq2, limma and rSeqRobust in negative-binomially distributed data. See Table 1 for descriptions.

DE (%)	Up (%)	edgeR-robust	DESeq2	limma-voom	rSeqRobust
30	50	**0.8865**	0.8862	0.8796	0.8845
		(0.0037)	(0.0042)	(0.0038)	(0.0037)
30	70	0.8678	0.8652	0.8616	**0.8721**
		(0.003)	(0.0033)	(0.0028)	**(0.0031)**
30	90	0.8428	0.8307	0.8316	**0.8551**
		(0.0018)	(0.0031)	(0.002)	**(0.0027)**
70	50	**0.8756**	0.8736	0.8687	0.8734
		(0.0032)	(0.0021)	(0.0034)	(0.0032)
70	70	0.8251	0.8071	0.8147	**0.8421**
		(0.0066)	(0.0044)	(0.007)	**(0.0029)**
70	90	0.6778	0.6536	0.6611	**0.756**
		(0.0072)	(0.0067)	(0.0074)	**(0.0075)**

This is hard to achieve with other robust methods such as the median or trimmed mean based normalization approaches (Anders and Huber, 2010; Robinson and Oshlack, 2010). This is attributed to the L0-penalized likelihood used by our model. Typically L0-penalized models are difficult to fit due to their non-convexity. However in our case it is easily manageable once we reduce the model fitting to a univariate or bivariate optimization problem. Adding an L0 penalty $p(\gamma)$ results in hard thresholding on γ, which has been shown to be a general case in She and Owen (2011). The hard thresholding facilitates the inference on indicators $\{\tau_i\}_{i=1}^m$ for differential gene expression.

The R program for generating the results in this paper is available at `http://www-personal.umich.edu/~jianghui/rseqrobust/`.

Appendix

Proof of Proposition 2.1. We have

$$\begin{aligned} f(\mu,\gamma,d) = & \sum_{i=1}^m \frac{1}{2\sigma_i^2}\left(\sum_{j=1}^{n_1}(x_{1ij}-\mu_i-d_{1j})^2+\sum_{s=2}^{S}\sum_{j=1}^{n_s}(x_{sij}-\mu_i-\gamma_{si}-d_{sj})^2\right) \\ & +\sum_{i=1}^m \alpha_i 1(\sum_{s=2}^{S}|\gamma_{si}|>0) \end{aligned}$$

Therefore

$$\frac{\partial f}{\partial d_{1j}} = \sum_{i=1}^m -\frac{1}{\sigma_i^2}(x_{1ij}-\mu_i-d_{1j}) = 0$$

which gives

$$d_{1j} = \frac{\sum_{i=1}^m \frac{1}{\sigma_i^2}(x_{1ij}-\mu_i)}{\sum_{i=1}^m \frac{1}{\sigma_i^2}}$$

Consequently

$$d_{1j}-d_{11} = \frac{\sum_{i=1}^m \frac{1}{\sigma_i^2}(x_{1ij}-x_{1i1})}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} = d'_{1j}$$

Since we fix $d_{11}=0$, we have $d_{1j}=d'_{1j}$. Similarly, we have

$$d_{sj}-d_{s1} = \frac{\sum_{i=1}^m \frac{1}{\sigma_i^2}(x_{sij}-x_{si1})}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} = d'_{sj}, s=2,\ldots,S$$

Denote d_{s1} as d_s, $s = 1, \ldots, S$, we have $d_1 = d_{11} = 0$ and $d_{sj} = d_s + d'_{sj}, s = 1, \ldots, S$. Now we have

$$f(\mu, \gamma, d) = \sum_{i=1}^{m} \left(\frac{1}{2\sigma_i^2} \sum_{j=1}^{n_1} (x_{1ij} - \mu_i - d'_{1j})^2 + \frac{1}{2\sigma_i^2} \sum_{s=2}^{S} \sum_{j=1}^{n_s} (x_{sij} - \mu_i - \gamma_{si} - d_s - d'_{sj})^2 + \alpha_i 1(\sum_{s=2}^{S} |\gamma_{si}| > 0) \right)$$

which can be written as

$$f(\mu, \gamma, d) = \sum_{i=1}^{m} h_i(\mu_i, \gamma_{2i}, \ldots, \gamma_{Si})$$

where the μ_i's and the γ_{si}'s can all be considered as functions of $(d_2, \ldots, d_S)$. Here we first work out these functions, i.e., solutions for the μ_i's and the γ_{si}'s when the d_s's are considered fixed. Fixing $d_2, \ldots, d_S$, f can be minimized by minimizing each h_i separately. When $\sum_{s=2}^{S} |\gamma_{si}| > 0$, h_i is easily minimized by

$$\mu_i = \frac{1}{n_1} \sum_{j=1}^{n_1} (x_{1ij} - d'_{1j}) = \mu'_{1i}$$

and

$$\gamma_{si} = \frac{1}{n_s} \sum_{j=1}^{n_s} (x_{sij} - \mu'_{1i} - d_s - d'_{sj}) = \frac{1}{n_s} \sum_{j=1}^{n_s} (x_{sij} - d'_{sj}) - \mu'_{1i} - d_s = \mu'_{si} - \mu'_{1i} - d_s$$

Similarly, when $\sum_{s=2}^{S} |\gamma_{si}| = 0$, i.e., $\gamma_{2i} = \cdots = \gamma_{Si} = 0$, we have

$$\mu_i = \frac{1}{\sum_{s=1}^{S} n_s} (n_1 \mu'_{1i} + \sum_{s=2}^{S} n_s(\mu'_{si} - d_s)) = \frac{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)}{n}$$

Changing from $\sum_{s=2}^{S} |\gamma_{si}| > 0$ to $\sum_{s=2}^{S} |\gamma_{si}| = 0$, the increase in the maximum achievable value of h_i is

$$h_i \left(\frac{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)}{n}, 0, \ldots, 0 \right) - h_i(\mu'_{1i}, \mu'_{2i} - \mu'_{1i} - d_2, \ldots, \mu'_{Si} - \mu'_{1i} - d_S)$$
$$= \frac{1}{2\sigma_i^2} \left\{ \sum_{s=1}^{S} n_s(\mu'_{si} - d_s)^2 - \frac{1}{n} \left[\sum_{s=1}^{S} (n_s(\mu'_{si} - d_s)) \right]^2 \right\} - \alpha_i$$

which is a quadratic function of $d_2, \ldots, d_S$. Therefore

$$\gamma_{si} = \begin{cases} 0 & \text{if } \frac{1}{2\sigma_i^2}\left\{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)^2 - \frac{1}{n}\left[\sum_{s=1}^{S}(n_s(\mu'_{si} - d_s))\right]^2\right\} < \alpha_i \\ \mu'_{si} - \mu'_{1i} - d_s & \text{otherswise} \end{cases}$$

Now we only need to solve for $d_2, \ldots, d_S$. We have

$$d_2, \ldots, d_S = \underset{d_2,\ldots,d_S}{\arg\min} \sum_{i=1}^{m} \min\left[h_i\left(\frac{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)}{n}, 0, \ldots, 0\right)\right.$$
$$\left., h_i(\mu'_{1i}, \mu'_{2i} - \mu'_{1i} - d_2, \ldots, \mu'_{Si} - \mu'_{1i} - d_S)\right]$$

which can be simplified as

$$d_2, \ldots, d_S = \underset{d_2,\ldots,d_S}{\arg\min} \sum_{i=1}^{m} \min\left(\frac{1}{2\sigma_i^2}\left\{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)^2 - \frac{1}{n}\left[\sum_{s=1}^{S}(n_s(\mu'_{si} - d_s))\right]^2\right\} - \alpha_i, 0\right)$$

since $h_i(\mu'_{1i}, \mu'_{2i} - \mu'_{1i} - d_2, \ldots, \mu'_{Si} - \mu'_{1i} - d_S)$ is actually not a function of $d_2, \ldots, d_S$. Therefore,

$$d_2, \ldots, d_S = \underset{d_2,\ldots,d_S}{\arg\min} \sum_{i=1}^{m} \min\left(\frac{1}{2\sigma_i^2}\left\{\sum_{s=1}^{S} n_s(\mu'_{si} - d_s)^2 - \frac{1}{n}\left[\sum_{s=1}^{S}(n_s(\mu'_{si} - d_s))\right]^2\right\}, \alpha_i\right)$$

□

References

S. Anders and W. Huber. Differential expression analysis for sequence count data. *Genome Biol*, 11(10):R106, 2010.

B. M. Bolstad, R. A. Irizarry, M. Astrand, and T. P. Speed. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics*, 19(2):185–193, Jan 2003.

M.-A. Dillies, A. Rau, J. Aubert, C. Hennequet-Antier, M. Jeanmougin, N. Servant, C. Keime, G. Marot, D. Castel, J. Estelle, G. Guernec, B. Jagla, L. Jouneau, D. Laloe, C. Le Gall, B. Schaeffer, L.ffer, S. Le

Crom, M. Guedj, F. Jaffrezic, and F. S. C. . A comprehensive evaluation of normalization methods for illumina high-throughput rna sequencing data analysis. *Brief Bioinform*, 14(6):671–683, Nov 2013.

J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96 (456):1348–1360, 2001.

J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.

T. J. Hastie, R. J. Tibshirani, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer, 2009.

H. Ji and X. S. Liu. Analyzing 'omics data using hierarchical models. *Nature biotechnology*, 28(4):337–340, 2010.

H. Ji and W. H. Wong. Tilemap: create chromosomal map of tiling array hybridizations. *Bioinformatics*, 21(18):3629–3636, 2005.

H. Jiang and W. H. Wong. Statistical inferences for isoform expression in rna-seq. *Bioinformatics*, 25(8):1026–1032, Apr 2009.

C. W. Law, Y. Chen, W. Shi, and G. K. Smyth. Voom: precision weights unlock linear model analysis tools for rna-seq read counts. *Genome Biol*, 15(2):R29, 2014.

B. Li, V. Ruotti, R. M. Stewart, J. A. Thomson, and C. N. Dewey. Rna-seq gene expression estimation with read mapping uncertainty. *Bioinformatics*, 26(4):493–500, Feb 2010.

J. Li, D. M. Witten, I. M. Johnstone, and R. Tibshirani. Normalization, testing, and false discovery rate estimation for rna-sequencing data. *Biostatistics*, 13(3):523–538, Jul 2012.

M. I. Love, W. Huber, and S. Anders. Moderated estimation of fold change and dispersion for rna-seq data with deseq2. *Genome Biology*, 15(12):550, 2014.

A. Mortazavi, B. A. Williams, K. McCue, L. Schaeffer, and B. Wold. Mapping and quantifying mammalian transcriptomes by rna-seq. *Nat Methods*, 5(7): 621–628, Jul 2008.

L. Pachter. Models for transcript quantification from RNA-Seq. *ArXiv e-prints*, Apr. 2011.

J. K. Pickrell, J. C. Marioni, A. A. Pai, J. F. Degner, B. E. Engelhardt, E. Nkadori, J.-B. Veyrieras, M. Stephens, Y. Gilad, and J. K. Pritchard. Understanding mechanisms underlying human gene expression variation with rna sequencing. *Nature*, 464(7289):768–772, 2010.

F. Rapaport, R. Khanin, Y. Liang, M. Pirun, A. Krek, P. Zumbo, C. E. Mason, N. D. Socci, and D. Betel. Comprehensive evaluation of differential gene expression analysis methods for rna-seq data. *Genome Biology*, 14 (9):R95, Sep 2013.

M. E. Ritchie, B. Phipson, D. Wu, Y. Hu, C. W. Law, W. Shi, and G. K. Smyth. limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Research*, 43(7):e47, Jan 2015.

M. D. Robinson and A. Oshlack. A scaling normalization method for differential expression analysis of rna-seq data. *Genome Biol*, 11(3):R25, 2010.

M. D. Robinson, D. J. McCarthy, and G. K. Smyth. edger: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26:139–140, Jan 2010.

J. Salzman, H. Jiang, and W. H. Wong. Statistical modeling of rna-seq data. *Statistical Science*, 26(1):62–83, 2011.

Y. She and A. B. Owen. Outlier detection using nonconvex penalized regression. *Journal of the American Statistical Association*, 106(494), 2011.

G. Smyth. Statistical applications in genetics and molecular biology. *Linear models and empirical Bayes methods for assessing differential expression in microarray experiments*, 2004.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

C. Trapnell, B. A. Williams, G. Pertea, A. Mortazavi, G. Kwan, M. J. van Baren, S. L. Salzberg, B. J. Wold, and L. Pachter. Transcript assembly and quantification by rna-seq reveals unannotated transcripts and isoform switching during cell differentiation. *Nat Biotechnol*, 28(5):511–515, May 2010.

G. P. Wagner, K. Kin, and V. J. Lynch. Measurement of mrna abundance using rna-seq data: Rpkm measure is inconsistent among samples. *Theory Biosci*, 131(4):281–285, Dec 2012.

Z. Wang, M. Gerstein, and M. Snyder. Rna-seq: a revolutionary tool for transcriptomics. *Nat Rev Genet*, 10(1):57–63, Jan 2009.

X. Zhou, H. Lindsay, and M. D. Robinson. Robustly detecting differential expression in rna sequencing data using observation weights. *Nucleic acids research*, 42(11):e91–e91, 2014.